A Probabilistic Approach Towards Exact-Repair Regeneration Codes

Mehran Elyasi
Department of ECE
University of Minnesota
melyasi@umn.edu

Soheil Mohajer
Department of ECE
University of Minnesota
soheil@umn.edu

arXiv:1510.02032v1 [cs.IT] 7 Oct 2015

***Abstract*—Regeneration codes with exact-repair property for distributed storage systems is studied in this paper. For exact-repair problem, the achievable points of (α, β) tradeoff match with the outer bound only for minimum storage regenerating (MSR), minimum bandwidth regenerating (MBR), and some specific values of n, k, and d. Such tradeoff is characterized in this work for general (n, k, k), (i.e., $k = d$) for some range of per-node storage (α) and repair-bandwidth (β). Rather than explicit code construction, achievability of these tradeoff points is shown by proving *existence* of exact-repair regeneration codes for any (n, k, k). More precisely, it is shown that an (n, k, k) system can be extended by adding a new node, which is randomly picked from some ensemble, and it is proved that, with high probability, the existing nodes together with the newly added one maintain properties of exact-repair regeneration codes. The new achievable region improves upon the existing code constructions. In particular, this result provides a complete tradeoff characterization for an $(n, 3, 3)$ distributed storage system for any value of n.**

I. Introduction

Distributed storage systems (DSS) are widely being used to provide reliability to storage technologies. Regeneration codes play a central role to manage data collection as well as system maintenance in DSS . An (n, k, d) regeneration-code encodes a file comprised of F symbols from a finite field $\mathbb{F}_q$ into n segments (nodes) $W_1, W_2, \ldots, W_n$, each including α symbols. Each data collector is able to recover the entire file by accessing any subset of nodes of size at least k. Moreover, whenever a node fails, it can be repaired by accessing d remaining nodes and downloading β symbols from each. The repair process can be performed in the functional or exact sense. In functional repair, a failed node will be replaced by another one, so that the resulting nodes maintain the data-recovery as well as node-repair properties. In exact repair (ER), however, the content of a failed node will be exactly replicated by the helpers.

It turns out that there is a fundamental tradeoff between minimum required values of α and β to store a given amount of data. Such tradeoff is derived for function-repair regeneration codes by Dimakis et al. [1], which is given by

$$F \leq \sum_{i=0}^{k-1} \min(\alpha, (d-i)\beta).$$

In practical applications, however, exact-repair is an appealing property, specially when it is desirable that the stored contents remain intact over time. In contrast to functional-repair, characterizing the optimal tradeoff between per-node storage and repair-bandwidth is a widely open problem for exact-repair regeneration codes.

This tradeoff is only characterized for very special cases. In particular, the optimum tradeoff of a $(4, 3, 3)$ system is characterized in [2] using a computer-aided approach, where it was shown that functional and exact repair tradeoffs are not identical. Moreover, the tradeoff is partially characterized in [3] for a $(5, 4, 4)$ system, which is extended to a complete characterization in [4]. For general (n, k, d)-DSS, families of outer bounds are developed, (e.g. [5], [6]), which can only partially characterize the tradeoff. Recently, the ER tradeoff is characterized for $(n = k + 1, k, d = k)$ systems, independently in [7]–[9], under the assumption of employing linear codes.

On the other hand, efficient ER regeneration codes are introduced for the minimum bandwidth regeneration (MBR) and minimum storage regeneration (MSR) points [10]–[12]. In spite of many other interesting proposed code construction (e.g. [13]–[15]), our knowledge about scalability of the system is very limited, for the interior of the tradeoff.

In this work, we study the scalability problem, and show that for any $n \geq k + 1$, the optimum tradeoff of an (n, k, k) ER system matches for some range of (α, β), with that of a $(k + 1, k, k)$ system. Unlike the standard approach in the literature where achievability of the tradeoff is shown by providing explicit code construction, we use a novel approach, based on random coding. Similar to Shannon's random coding argument, we show that an existing code can be extended by appending new randomly generated nodes, which with high probability maintain the exact-repair property. To the best of our knowledge, this is the first result which suggests that the ER tradeoff may only depend on (k, d, α, β) and not the number of nodes n for the interior of the region (similar to the functional repair case). In particular, our approach yields in a complete ER tradeoff characterization for an $(n, 3, 3)$ DSS.

The rest of this paper is organized as follows. We first formally present the results in Section II. In Section III we characterize a set of conditions a newly added node should satisfy in order to maintain ER property, and in Section IV

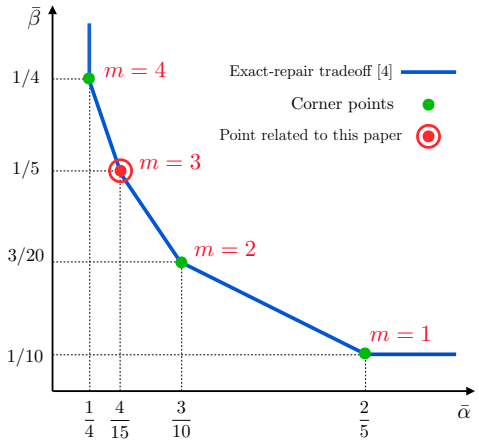

Fig. 1. The exact-repair tradeoff for $(5,4,4)$-DSS.

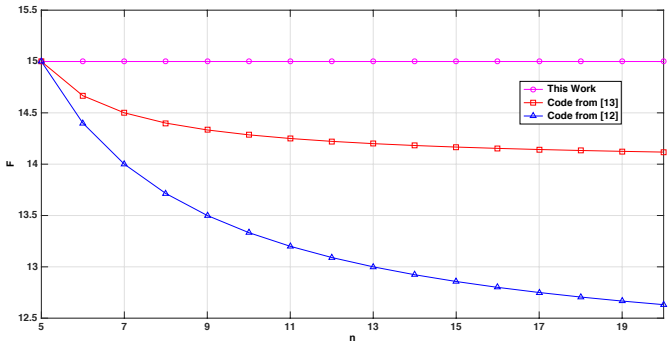

Fig. 2. ER capacity of $(n,4,4)$-DSS as a function n for $(\alpha,\beta)=(4,3)$.

we show that w.h.p. a randomly chosen node satisfies these properties. Built on the tools developed in Sections III and IV, the proof of the main results are presented in Section V.

II. The Model and Main Result

An exact-repair regenerating distributed storage system with parameters (n,k,d) and (α,β) consists of n storage nodes, each with storage capacity α symbols. A file including F symbols from some finite field $\mathbb{F}_q$ is encoded into n pieces of information, in distributed manner, and each piece is stored on one of the storage nodes. We associate each node with an index in $I_n \triangleq \{1,2,\ldots,n\}$, and denote by W_i the content of node i, for $i \in I_n$. *Data recovery* property implies that the original file can be recovered from any subset of nodes $\mathcal{A} \subseteq I_n$, provided that $|\mathcal{A}| \geq k$. Moreover, if any node $x \in I_n$ fails, its content $\mathcal{W}_x$ can be duplicated by receiving (at most) β symbols from each node in a set $\mathcal{A}$ (called *helper nodes*), for every $\mathcal{A} \subseteq I_n \setminus \{x\}$ with $|\mathcal{A}| \geq d$. For a given file size F, characterizing the optimum tradeoff between α and β satisfying the aforementioned constraints is a challenging open problem for general (n,k,d)-DSS.

Recently, the optimum tradeoff between the per-node storage and repair-bandwidth of ER regeneration codes is characterized in [7]–[9] for an $(n=k+1,k,d=k)$-DSS, under the limitation of employing linear codes. It is shown that the optimum tradeoff for a $(k+1,k,k)$-ER linear DSS is a piecewise-linear function, with k corner point. The m-th corner points of this region is given by

$$\bar{\alpha} = \frac{m+1}{m(k+1)} \qquad \text{and} \qquad \bar{\beta} = \frac{m+1}{k(k+1)}, \tag{1}$$

for $m=1,2,\ldots,k$.

In particular, it was shown that a class of codes proposed in [13] and [14] achieve this optimum tradeoff for any $(k+1,k,k)$. Moreover, the linearity constraint is relaxed for $k=3$ and $k=4$ in [2] and [4], respectively, where the optimality of the tradeoff is proved using general information-theoretic arguments. This tradeoff is shown for $k=4$ in Figure 1.

The two extreme points in (1), namely $m=1$ and $m=k$, correspond to the MBR and MSR points, respectively. Code constructions for MBR point are provided in [10] for general $(n,k,d=k)$. Moreover, it is shown that these parameters are *asymptotically* achievable for the MSR point [11], [12]. However, it is not clear whether the interior of the tradeoff in (1) can be achieved when the number of storage nodes is larger than $k+1$, that is $n > k+1$.

In this work we provide a partial answer to this question in a positive way: *the corner points in* (1) *on the optimum tradeoff associated with* $m=k-1$ *is achievable*. We use a random coding strategy and probabilistic argument, to show that for any set of $n \geq k+1$ storage nodes forming (n,k,k) exact-repair regeneration code, one can always *find* a new node, such that the new node together with the existing n nodes form an $(n+1,k,k)$ exact-repair regeneration code. Instead of constructing a new node, we rather pick it from an ensemble of nodes, and show that a random node preserves the desired properties with high probability, provided that the underlying field size is large enough.

The main result of this work is formally stated in the following theorem.

Theorem 1. *For any $(n,k,d=k)$ distribued storage system with $n > k$, there exist some large enough q, and exact repair regeneration codes over $\mathbb{F}_q$ with normalized per-node capacity and repair-bandwidth*

$$(\bar{\alpha},\bar{\beta}) = \left(\frac{m+1}{m(k+1)}, \frac{m+1}{k(k+1)}\right) = \left(\frac{k}{k^2-1}, \frac{k-1}{k^2-1}\right).$$

The performance of the code proposed in this paper for $k=4$ is compared to that of existing ones in the literature in Figure 2. As it is shown, for both codes proposed in [13] and [14] the maximum file size can be stored in an ER regeneration system for given $(\alpha,\beta)=(k,k-1)$ decreases with n, whereas number of nodes is irrelevant to the file size in the new proposed code.

Note that existence of such ER regeneration codes provides a partial characterization for the optimum tradeoff of an (n,k,k)-DSS as follows.

Corollary 1. *The optimum tradeoff between the per-node capacity and repair-bandwidth of any distributed storage system with parameters (n,k,k) is given by*

$$\bar{\beta} + (k-1)\bar{\alpha} \geq 1,$$

for $(k-1)/k \leq \bar{\alpha}/\bar{\beta} \leq 1$.

As another consequence, Theorem 1 suffices to characterize the entire optimum tradeoff for an $(n,3,3)$-DSS.

Corollary 2. *The optimum tradeoff between the per-node capacity and repair-bandwidth of any exact-repair* $(n,k=3,d=3)$ *distributed storage system is given by*

$$3\bar{\alpha} \geq 1, \qquad 2\bar{\alpha}+\bar{\beta} \geq 1, \qquad 4\bar{\alpha}+6\bar{\beta} \geq 3, \qquad 6\bar{\beta} \geq 1.$$

The proof of these results can be found in Section V.

III. Code Structure and Well-Aligned Node

We start by redefining exact-repair regeneration codes in a linear framework.

A. Linear Codes

It is easier to analyze linear regeneration codes in the context of vector spaces over a finite field. In this context, we denote subspaces by script letters, e.g. $\mathcal{W}, \mathcal{S}$, and with slightly abuse of notation, use $\subseteq$ to denote subspace relationship. Moreover, we use $\langle\cdot\rangle$ to denote the span of a set of vectors. Furthermore, we adopt the notation used in [7] as follows:

A linear exact-repair regeneration code with parameters (α,β,F) for an (n,k,d) distributed storage system can be defined as

- An F-dimensional vector space in $\mathbb{F}_q$ for some q, which we denote by $\mathcal{F}$. The F data symbols stored in the DSS form a basis for this vector space.
- There are a total of α vectors from $\mathcal{F}$ stored in node i, for $i \in I_n$. These vectors span a subspace $\mathcal{W}_i \subseteq \mathcal{F}$, with $\dim(\mathcal{W}_i) \leq \alpha$.
- *Data recovery*: By accessing any subset of k nodes, the entire vector space $\mathcal{F}$ can be spanned. In other words,
$$\mathcal{F} = \sum_{j\in\mathcal{A}} \mathcal{W}_j \qquad \forall \mathcal{A} \subseteq I_n, |\mathcal{A}| = k.$$
- *Node repair*: In the case of failure of node x, its content can be spanned by summation of d *helping* vector spaces, each of dimensional β and coming from one other one. More precisely, for every subset of nodes $\mathcal{A} \subseteq I_n \setminus \{x\}$ with $|\mathcal{A}| = d$, we have[1]
$$\mathcal{W}_x \subseteq \sum_{j\in\mathcal{A}} \mathcal{S}_j^x[\mathcal{A}]$$
where $\mathcal{S}_j^x[\mathcal{A}] \subseteq \mathcal{W}_j$ and $\dim\left(\mathcal{S}_j^x[\mathcal{A}]\right) \leq \beta$.

The following proposition lists some structural properties of optimum linear ER regeneration codes.

Proposition 1. *Let* $\mathcal{C}$ *be any* (n,k,k) *ER regeneration code operating at* $(\alpha,\beta,F) = (k, k-1, k^2-1)$. *Consider the repair process of node* $x \in I_n$ *via nodes in* $\mathcal{A}$, *where* $\mathcal{A} \subseteq I_n \setminus \{x\}$, *and* $|\mathcal{A}| = k$. *Then vector space of each node* $j \in \mathcal{A}$ *can be partitioned into two subspaces* $\mathcal{S}_j^x[\mathcal{A}]$ *and* $\mathcal{T}_j^x[\mathcal{A}]$ *such that*

(i) $\mathcal{S}_j^x[\mathcal{A}]$ *is the subspace node* j *sends to repair node* x, *and* $\dim\left(\mathcal{S}_j^x[\mathcal{A}]\right) = \beta = (k-1)$;

(ii) $\mathcal{T}_j^x[\mathcal{A}] = \left\langle t_j^x[\mathcal{A}]\right\rangle$ *is a one-dimensional subspace spanned by* $t_j^x[\mathcal{A}] \in \mathcal{W}_j$;

(iii) $\mathcal{W}_j = \mathcal{S}_j^x[\mathcal{A}] \oplus \mathcal{T}_j^x[\mathcal{A}]$;

(iv) $\sum_{j\in\mathcal{A}} t_j^x[\mathcal{A}] = 0$;

(v) $\mathcal{T}^x[\mathcal{A}] \triangleq \sum_{j\in\mathcal{A}} \mathcal{T}_j^x[\mathcal{A}]$ *is a* $(k-1)$*-dimensional subspace of* $\mathcal{F}$;

(vi) $\mathcal{T}^x[\mathcal{A}]$ *is spanned by* $\{t_j^x[\mathcal{A}] : j \in \mathcal{B}\}$, *for every* $\mathcal{B} \subseteq \mathcal{A}$ *with* $|\mathcal{B}| = k-1$.

Partitioning of node spaces is given for a $(4,3,3)$ system in Figure 5, which can be useful to follow the statements above. This proposition plays a central role in our arguments. The proof of the proposition can be found in Appendix A.

B. Well-Aligned nodes

Our approach to achieve (α,β) for an (n,k,k)-DSS is recursive, i.e., we start with an existing $(k+1,k,k)$ ER code, and in each step append one new node to the system. Such new node should satisfy a set of conditions so that system maintain the data-recovery and exact-repair properties. It will be later shown that such conditions are fulfilled by any node satisfy the following definition.

Definition 1. *Let* $\mathcal{C}$ *be an* (n,k,k) *ER regeneration code with parameters* $(\alpha,\beta,F) = (k,k-1,k^2-1)$. *Fix* $\mathcal{A} \subseteq I_n$ *with* $|\mathcal{A}| = k$, *and let* $x \notin \mathcal{A}$ *be a node index. An* α*-dimensional subspace (node)* $\mathcal{W}_\star$ *is called* well-aligned with respect to a pair $(\mathcal{A};x)$ *if it can be spanned by a set of basis vectors* $\mathscr{B} = \{w_{\star(1)}, w_{\star(2)}, \ldots, w_{\star(k)}\}$, *such that [see Figure 3]*

(1) $w_{\star(i)} = \sum_{j\in\mathcal{A}} s_{j(i)} + \tau_{(i)}$ *for every* $i = 1,\ldots,k$, *where* $s_{j(i)} \in \mathcal{S}_j^x[\mathcal{A}]$, *and* $\tau_{(i)} \in \mathcal{T}^x[\mathcal{A}]$;

(2) *for each* $j \in \mathcal{A}$, *there exists* $i_j \in I_k$ *such that* $s_{j(i_j)} = \mathbf{0}$;

(3) *for every* $j \in \mathcal{A}$, *set of vectors*
$$\{s_{j(1)}, s_{j(2)}, , \ldots, s_{j(k)}\} \setminus \{s_{j(i_j)}\}$$
are $(k-1)$ *linearly independent vectors from* $\mathcal{S}_j^x[\mathcal{A}]$.

The next proposition shows that a well-aligned node w.r.t. $(\mathcal{A},x)$ satisfies some desired properties for an exact repair regeneration code.

Proposition 2. *Let* $\mathcal{C}$ *be an* (n,k,k) *ER regeneration code with parameters* $(\alpha,\beta,F) = (k,k-1,k^2-1)$. *Consider* $\mathcal{A}$ *to be a subset of* k *nodes, and* $x \in I_n \setminus \mathcal{A}$. *Let* $\mathcal{W}_\star$ *be a node that is well-aligned w.r.t.* $(\mathcal{A};x)$. *Then* $\mathcal{C}^\star = \{\mathcal{W}_i; i \in \mathcal{A}\} \cup \{\mathcal{W}_\star\}$ *form a* $(k+1,k,k)$ *ER regeneration code with the same parameters* (α,β,F).

Proof of Proposition 2: In order to prove this proposition we need to show data-recovery and node repair properties of $\mathcal{C}^*$.

[1] Note that the repair sent by node j to x potentially depends on the other helpers. This is captured by $[\mathcal{A}]$ in our notation $\mathcal{S}_j^x[\mathcal{A}]$. However, we may drop $[\mathcal{A}]$, whenever either $\mathcal{A}$ is unique, or dependency of $\mathcal{A}$ is clear from the context.

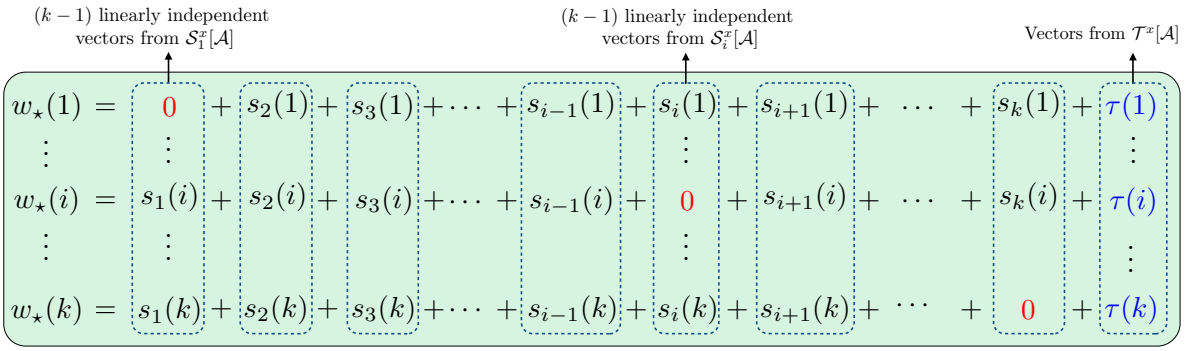
$(k-1)$ linearly independent vectors from $\mathcal{S}_1^x[\mathcal{A}]$
$(k-1)$ linearly independent vectors from $\mathcal{S}_i^x[\mathcal{A}]$
Vectors from $\mathcal{T}^x[\mathcal{A}]$
$w_\star(1) = 0 + s_2(1) + s_3(1) + \cdots + s_{i-1}(1) + s_i(1) + s_{i+1}(1) + \cdots + s_k(1) + \tau(1)$
$w_\star(i) = s_1(i) + s_2(i) + s_3(i) + \cdots + s_{i-1}(i) + 0 + s_{i+1}(i) + \cdots + s_k(i) + \tau(i)$
$w_\star(k) = s_1(k) + s_2(k) + s_3(k) + \cdots + s_{i-1}(k) + s_i(k) + s_{i+1}(k) + \cdots + 0 + \tau(k)$

Fig. 3. Structure of a well-aligned node with respect to $\mathcal{A} = \{1, 2, \ldots, k\}$.

Node repair: Note that $\mathcal{C}^\star$ has only $(k+1)$ nodes, and once the failed node $x \in \mathcal{A} \cup \{\star\}$ is chosen, the helper nodes are uniquely determined. To simplify the presentation, w.o.l.g. we may assume $\mathcal{A} = \{1, 2, \ldots, k\}$ and $i_j = j$ for every $j \in \mathcal{A}$, i.e., $s_{j(j)} = \mathbf{0}$. We can identify the following two cases.

1) *Repair of $\mathcal{W}_\star$ using nodes in $\mathcal{A}$:* Note that $\mathcal{W}_\star = \langle \mathscr{B} \rangle$. For this repair process, we reconstruct each vector in $\mathscr{B}$ using repair vectors sent from nodes in $\mathcal{A}$. Recall by part (vi) of Proposition 1 that since $\tau(i) \in \mathcal{T}^x$, it can be represented as

$$\tau(i) = \sum\nolimits_{j \in \mathcal{A}\setminus\{i\}} \theta_{i,j} t_j^x [\mathcal{A}]$$

for every $i \in \mathcal{A}$. Now, we define the repair subspace $\mathcal{S}_j^\star [\mathcal{A}]$ sent by node $j \in \mathcal{A}$ to node $\star$ by

$$\mathcal{S}_j^\star [\mathcal{A}] = \left\langle \left\{ s_{j(i)} + \theta_{i,j} t_j^x [\mathcal{A}] : j = 1, 2, \ldots, k, j \neq i \right\} \right\rangle.$$

It is clear that $\mathcal{S}_j^\star [\mathcal{A}]$ is spanned by only $(k-1)$ vectors, and hence its dimension does not exceed β. Now, vector $w_{\star(i)} \in \mathscr{B}$ can be reconstructed via

$$\begin{aligned} w_{\star(i)} &= \sum_{j \in \mathcal{A}\setminus\{i\}} s_{j(i)} + \tau(i) \\ &= \sum_{j \in \mathcal{A}\setminus\{i\}} \left(s_{j(i)} + \theta_{i,j} t_j^x [\mathcal{A}] \right) \in \sum_{j \in \mathcal{A}\setminus\{i\}} \mathcal{S}_j^* [\mathcal{A}]. \end{aligned}$$

2) *Repair of node $\ell \in \mathcal{A}$ using nodes in $\mathcal{B} \triangleq \{\star\} \cup (\mathcal{A} \setminus \{\ell\})$:* This repair process is more technical. However, the illustration in Figure 4 can be helpful to follow the proof. Recall that $\mathcal{W}_\ell = \mathcal{S}_\ell^x [\mathcal{A}] \oplus \langle t_j^x [\mathcal{A}] \rangle$. In order to repair $\mathcal{W}_\ell$, we rebuild $\mathcal{S}_\ell^x [\mathcal{A}]$ and $\langle t_j^x [\mathcal{A}] \rangle$, separately. To reconstruct $\mathcal{S}_\ell^x [\mathcal{A}]$, node $\mathcal{W}_\star$ sends a subspace

$$\mathcal{S}_\star^\ell [\mathcal{B}] \triangleq \langle \{ w_{\star(i)} : i \neq \ell \} \rangle.$$

Note that $\mathcal{S}_\star^\ell [\mathcal{B}]$ is spanned by $(k-1)$ linearly independent vectors, thus $\dim\left(\mathcal{S}_\star^\ell [\mathcal{B}]\right) = k - 1 = \beta$. Moreover, each vector in $\{w_{\star(i)} : i \neq \ell\}$ can be written as

$$w_{\star(i)} = s_{\ell(i)} + \sum_{j \in \mathcal{A}\setminus\{i,\ell\}} s_{j(i)} + \tau(i), \qquad (2)$$

which is indeed a vector from $\mathcal{S}_\ell^x [\mathcal{A}]$ that is corrupted by some interference.

The repair data sent by other nodes in $\mathcal{A} \setminus \{\ell\}$ play two important roles: (i) cancel the interference in $w_{\star(i)}$'s, and (ii) recover the vector space $\langle t_\ell^x [\mathcal{A}] \rangle$. Let us define

$$\mathcal{S}_j^\ell [\mathcal{B}] = \langle \{ s_{j(i)} : i \in \mathcal{A} \setminus \{j, \ell\} \} \rangle \oplus \langle t_j^x [\mathcal{A}] \rangle.$$

First, it is clear that $\mathcal{S}_j^\ell [\mathcal{B}] \subseteq \mathcal{W}_j$. Moreover since $|\{ s_{j(i)} : i \in \mathcal{A} \setminus \{j, \ell\} \}| = k - 2$, we have $\dim\left(\mathcal{S}_j^\ell [\mathcal{B}]\right) \leq (k-2) + 1 = k - 1 = \beta$. Now, from part (vi) of Proposition 1 node ℓ can first recover $\mathcal{T}^x [\mathcal{A}]$ from the $(k-1)$ received vectors $\{t_j^x [\mathcal{A}] : j \in \mathcal{A}, j \neq \ell\}$. Once $\mathcal{T}^x [\mathcal{A}]$ is rebuilt, the subspace $\mathcal{T}_\ell^x [\mathcal{A}] \subseteq \mathcal{T}^x [\mathcal{A}]$ can be also recovered.
Furthermore, since $\tau(i) \in \mathcal{T}^x [\mathcal{A}]$, they can be all reconstructed and canceled from $w_{\star(i)}$ in (2). The remaining interference in $w_{\star(i)}$, given by $\sum_{j \in \mathcal{A}\setminus\{i,\ell\}} s_{j(i)}$, can be canceled since $s_{j(i)} \in \mathcal{S}_j^\ell [\mathcal{B}]$ for every $j \in \mathcal{A} \setminus \{i, \ell\}$. By removing all interference, vectors

$$\left\{ s_{\ell(1)}, \ldots, s_{\ell(\ell-1)}, s_{\ell(\ell+1)}, \ldots, s_{\ell(k)} \right\}$$

can be reconstructed at the failed node. Then, since these vectors are linearly independent (by part (3) of Definition 1), they can completeley span $\mathcal{S}_\ell^x [\mathcal{A}]$, which together with $\mathcal{T}_\ell^x [\mathcal{A}]$, can span $\mathcal{W}_\ell$.

Data recovery: Recovering file from nodes in $\mathcal{A}$ is clear, since nodes in $\mathcal{A}$ were already a part of an (n, k, k) ER code. Consider a set $\mathcal{B} = \{*\} \cup \mathcal{A} \setminus \{\ell\}$. The argument used to proof repairability of $\mathcal{W}_\ell$ shows that

$$\mathcal{S}_\ell^x [\mathcal{A}] \subseteq \sum_{i \in \mathcal{B}} \mathcal{S}_i^\ell [\mathcal{B}] \subseteq \sum_{i \in \mathcal{B}} \mathcal{W}_i.$$

Moreover, part (vi) of Proposition 1 implies

$$\mathcal{T}_\ell^x [\mathcal{A}] \subseteq \mathcal{T}^x [\mathcal{A}] = \sum_{i \in \mathcal{A}\setminus\{\ell\}} \mathcal{T}_i^\ell [\mathcal{B}] \subseteq \sum_{i \in \mathcal{A}\setminus\{\ell\}} \mathcal{W}_i.$$

Hence, we have

$$\begin{aligned} \mathcal{F} &= \mathcal{W}_\ell + \sum_{i \in \mathcal{A}\setminus\{\ell\}} \mathcal{W}_i \subseteq (\mathcal{S}_\ell^x [\mathcal{A}] + \mathcal{T}_\ell^x [\mathcal{A}]) + \sum_{i \in \mathcal{A}\setminus\{\ell\}} \mathcal{W}_i \\ &\subseteq \mathcal{W}_\star + \sum_{i \in \mathcal{A}\setminus\{\ell\}} \mathcal{W}_i = \sum_{i \in \mathcal{B}} \mathcal{W}_i, \end{aligned}$$

which implies that $\mathcal{F}$ can be recovered from nodes in $\mathcal{B}$. This completes the proof of the data-recovery property. ■

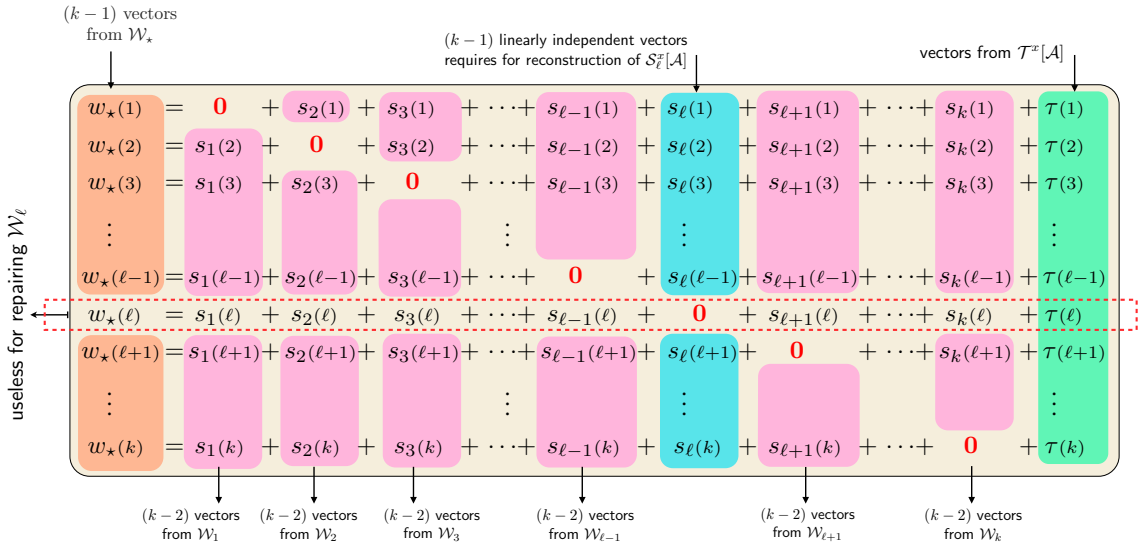
$(k-1)$ vectors from $\mathcal{W}_\star$
$(k-1)$ linearly independent vectors requires for reconstruction of $\mathcal{S}_\ell^x[\mathcal{A}]$
vectors from $\mathcal{T}^x[\mathcal{A}]$
useless for repairing $\mathcal{W}_\ell$
$(k-2)$ vectors from $\mathcal{W}_1$
$(k-2)$ vectors from $\mathcal{W}_2$
$(k-2)$ vectors from $\mathcal{W}_3$
$(k-2)$ vectors from $\mathcal{W}_{\ell-1}$
$(k-2)$ vectors from $\mathcal{W}_{\ell+1}$
$(k-2)$ vectors from $\mathcal{W}_k$

Fig. 4. Repair process of node ℓ via $\{\star\} \cup \{1, 2, \ldots, \ell-1, \ell+1, \ldots, k\}$.

Remark 1. *Note that a subspace $\mathcal{W}_\star$ can be well aligned w.r.t. $(\mathcal{A}; x)$ and not be well-aligned w.r.t. $(\mathcal{A}; y)$ for $x \neq y$. However, nodes in $\mathcal{A}$ together with $\mathcal{W}_\star$ form a $(k+1, k, k)$ ER regeneration code as long as there exist at least one x such that $\mathcal{W}_\star$ is well-aligned w.r.t. $(\mathcal{A}; x)$.*

IV. Probabilistic Method for Appending New Nodes

Consider an (n, k, k) ER regeneration code, $\mathcal{C} = \{\mathcal{W}_i : i \in I_n\}$. Our goal is to append a new node $\mathcal{W}_\star$ to this system, so that $\tilde{\mathcal{C}} = \mathcal{C} \cup \{\mathcal{W}_\star\}$ maintain properties of ER regeneration codes. Such a code will be an $(n+1, k, k)$ ER code.

It is easy to see that the necessary and sufficient condition for a new node $\mathcal{W}_\star$ to be feasible to be added to $\tilde{\mathcal{C}}$ is the following: $\tilde{\mathcal{C}}$ is an $(n+1, k, k)$ ER regeneration code if and only if $\mathcal{C}_A^\star = \{\mathcal{W}_i : i \in \mathcal{A}\} \cup \{\mathcal{W}_\star\}$ form a $(k+1, k, k)$ ER regeneration code for every $\mathcal{A} \subseteq I_n$ with $|\mathcal{A}| = k$.

Furthermore, Proposition 2 shows that in order to $\mathcal{C}_A^\star = \{\mathcal{W}_i : i \in \mathcal{A}\} \cup \{\mathcal{W}_\star\}$ be a $(k+1, k, k)$ ER code, it suffices that $\mathcal{W}_\star$ be well-aligned w.r.t. $(\mathcal{A}, x)$ for some $x \in I_n \setminus \mathcal{A}$.

In this section we will show that a randomly chosen α-dimensional subspace $\mathcal{W}_\star$ is well-aligned w.r.t. every choice of $(\mathcal{A}; x)$ (with $|\mathcal{A}| = k$) for some $x \in \mathcal{A}$, with high probability, and hence can be add to the current code with n nodes. More formally, we prove the following proposition.

Proposition 3. *Let $\mathcal{C} = \{\mathcal{W}_i : i \in I_n\}$ be an (n, k, k) ER regeneration code storing a data space $\mathcal{F}$ over $\mathbb{F}_q$, and $\mathcal{W}_\star$ be an α-dimensional subspace of $\mathcal{F}$ drawn uniformly at random. Then for any $\epsilon > 0$, there exists a large enough q such that*

$$\mathbb{P}[\mathcal{C} \cup \{\mathcal{W}_\star\} \text{ is an } (n+1, k, k) \text{ ER code}] > 1 - \epsilon.$$

The rest of this section is dedicated to prove Proposition 3. Let $E(\mathcal{A}; x)$ be the event that a random $\mathcal{W}_\star$ is well-aligned w.r.t. $(\mathcal{A}; x)$ for $\mathcal{A} \subseteq I_n$ and $x \notin \mathcal{A}$. Moreover, define

$$E(\mathcal{A}) = \bigcup_{x \in I_n \setminus \mathcal{A}} E(\mathcal{A}; x),$$

which is the event of existence of some x for which $\mathcal{W}_\star$ is well-aligned w.r.t. $(\mathcal{A}, x)$.

Now let $\mathcal{W}_\star$ be a subspace of $\mathcal{F}$ picked uniformly at random among all α-dimensional subspaces. Then we have

$$\begin{aligned}
\mathbb{P}\left[\bigcap_{\mathcal{A} \subset I_n, |\mathcal{A}|=k} E(\mathcal{A})\right] &= 1 - \mathbb{P}\left[\bigcup_{\mathcal{A} \subset I_n, |\mathcal{A}|=k} E^c(\mathcal{A})\right] \\
&\geq 1 - \sum_{\mathcal{A} \subset I_n, |\mathcal{A}|=k} \mathbb{P}\left[E^c(\mathcal{A})\right] \\
&= 1 - \sum_{\mathcal{A} \subset I_n, |\mathcal{A}|=k} (1 - \mathbb{P}\left[E(\mathcal{A})\right]) \\
&= 1 - \binom{n}{k} (1 - \mathbb{P}\left[E(\mathcal{A}_0)\right]), \qquad (3)
\end{aligned}$$

where $\mathcal{A}_0 = \{1, 2, \ldots, k\}$ is a fixed subset of I_n. here, the last equality is due to the symmetric structure of nodes, which implies $\mathbb{P}[E(\mathcal{A})]$ does not depend on the realization of $\mathcal{A}$.

Next, let x_0 be a fixed node in $I_n \setminus \mathcal{A}_0$. We have $E(\mathcal{A}_0; x_0) \subseteq E(\mathcal{A}_0)$. Hence,

$$\mathbb{P}\left[E(\mathcal{A}_0)\right]) \geq \mathbb{P}\left[E(\mathcal{A}_0; x_0)\right]),$$

Therefore, we can further lower bound RHS of (3) by

$$\mathbb{P}\left[\bigcap_{\mathcal{A} \subset I_n, |\mathcal{A}|=k} E(\mathcal{A})\right] \geq 1 - \binom{n}{k} (1 - \mathbb{P}\left[E(\mathcal{A}_0; x_0)\right]). \qquad (4)$$

Thus, in order to show $\mathbb{P}\left[\bigcap_{\mathcal{A} \subset I_n, |\mathcal{A}|=k} E(\mathcal{A})\right] \xrightarrow{q \to \infty} 1$, it suffices to prove $\mathbb{P}\left[E(\mathcal{A}_0; x_0)\right] \to 0$ as q grows. To this end, we can evaluate $\mathbb{P}\left[E(\mathcal{A}_0)\right]$ by

$$\mathbb{P}\left[E(\mathcal{A}_0; x_0)\right] = \frac{\#\text{ well-aligned subspaces w.r.t } (A_0; x_0)}{\#\ \alpha\text{-dimesnional subspaces of } \mathcal{F}}. \qquad (5)$$

It is well-known (e.g. see [16]) that the denominator in (5) for $\dim(\mathcal{F}) = F$ is given by

$$\frac{\prod_{h=0}^{\alpha-1}(q^F - q^h)}{\prod_{h=0}^{\alpha-1}(q^\alpha - q^h)} = \frac{\prod_{h=0}^{k-1}(q^{k^2-1} - q^h)}{\prod_{h=0}^{k-1}(q^k - q^h)} \tag{6}$$

In order to compute the nominator, we must count the total number of well-aligned nodes w.r.t. a fixed pair $(\mathcal{A}_0, x_0)$. Note that for each $i \in \mathcal{A}$, we need to pick a total of $(k-1)$ vectors, namely $\{s_{i(j)} : j \in \mathcal{A} \setminus \{i\}\}$, from a $(k-1)$-dimensional space $\mathcal{S}_i^{x_0}[\mathcal{A}]$ (see the i-th column in Figure 3). For a given i, this can be done in $\prod_{h=0}^{k-2}(q^{k-1} - q^h)$ ways.

On the other hand, vectors $\tau(j)$ in the last column are picked arbitrarily and independent of each other from an $(k-1)$-dimensional space $\mathcal{T}^x[\mathcal{A}]$. Therefore there are a total q^{k-1} choices for each $\tau(i)$.

Hence, the number of choices for the basis set of $\mathcal{W}_\star$ is

$$\prod_{i=1}^{k}\left[\prod_{h=0}^{k-2}(q^{k-1} - q^h)\right] \cdot \prod_{j=1}^{k} q^{k-1} = \left[\prod_{h=0}^{k-2}(q^{k-1} - q^h)\right]^k q^{k(k-1)}.$$

Finally note that once vectors $s_{j(i)}$ and $\tau(j)$ are fixed for $i, j \in I_k$, each basis vector $w_\star(j)$ in Figure 3 can be scaled by any non-zero $\xi \in \mathbb{F}_q$, while the resulting vector space is preserved. Considering this fact for k basis vectors of $\mathcal{W}_\star$, the nominator of (5) can be evaluated by

$$\frac{\left[\prod_{h=0}^{k-2}(q^{k-1} - q^h)\right]^k \cdot q^{k(k-1)}}{q^k}. \tag{7}$$

Replacing (6) and (7) in (5), we get

$$\begin{aligned}\mathbb{P}[E(\mathcal{A}_0; x_0)] &= \frac{\left[\prod_{h=0}^{k-2}(q^{k-1} - q^h)\right]^k \cdot q^{k(k-1)} \cdot \prod_{h=0}^{k-1}(q^k - q^h)}{\left[\prod_{h=0}^{k-1}(q^{k^2-1} - q^h)\right] \cdot q^k} \\ &= \frac{q^{k^3}(1 - o(1))}{q^{k^3}(1 - o(1))} = 1 - o(1),\end{aligned} \tag{8}$$

where $o(1)$ vanishes as $q \to \infty$. This completes the proof.

V. Proof of the Main Results

We have developed all the techniques need to prove the main results of Section II in the previous sections.

Proof of Theorem 1: We present the proof of the theorem using a recursive argument. We first note that for $(k+1, k, k)$ DSS, exact-repair regeneration codes with parameters $(\alpha, \beta, F) = (k, k-1, k^2-1)$ are introduced independently in [13] and [14]. We use them as the starting point of the recursive argument.

Moreover, Proposition 3 implies that, for large enough q, a randomly sampled α-dimensional subspace can be appended to an (n, k, k) system to form an $(n+1, k, k)$ exact-repair regeneration code. Hence, we can start with $n = k+1$, and repeat picking random new nodes $\mathcal{W}_\star$ and checking whether they satisfy the desired properties. By repeating this procedure we can get as many number of nodes needed while the entire system preserves the exact-repair property. ∎

Proof of Corollary 1: The optimality of this tradeoff can be simply seen using a cut-set argument

$$\begin{aligned}F &= H(W_1, \ldots, W_k) \\ &\le H(W_1, \ldots, W_{k-1}) + H(W_k | W_1, \ldots, W_{k-1}) \\ &\le (k-1)\alpha + \beta,\end{aligned}$$

which implies $\bar{\beta} \ge 1 - (k-1)\bar{\alpha}$. Moreover, for $(k-1)/k \le \bar{\alpha}/\bar{\beta} \le 1$, the two extreme points of this bound are given by

$$(\bar{\alpha}, \bar{\beta}) = \left(\frac{1}{k}, \frac{1}{k}\right) \qquad \text{and} \qquad (\bar{\alpha}, \bar{\beta}) = \left(\frac{k}{k^2-1}, \frac{k-1}{k^2-1}\right).$$

Note that the first point is MSR, for which achievability is known for arbitrary value of n [11], [12]. Achievability of the second point is proved in Theorem 1. This completes the proof. ∎

Proof of Corollary 2: Note that since an $(n, 3, 3)$ DSS includes a $(4, 3, 3)$ DSS, its exact-repair tradeoff cannot be lower than that of a $(4, 3, 3)$ DSS. On the other hand, the exact repair tradeoff for a $(4, 3, 3)$ system is characterized by Tian [2] as given in the corollary.

In order to show that this tradeoff is indeed achivable, we can focus on the corner points, since the intermediate points can be achieved by space-sharing. There are three corner points for this region, namely, MBR, MSR, and one given by $(\bar{\alpha}, \bar{\beta}) = \left(\frac{k}{k^2-1}, \frac{k-1}{k^2-1}\right)$. Achievability of the first two points is known for arbitrary n [10]–[12]. For the middle point, however, Theorem 1 guarantees existence of n nodes maintaining data-recovery and exact-repair prperties, for large enough q. Hence, the entire boundary of the tradeoff is achievable for any value of $n \ge k+1$. ∎

Appendix A
Proof of proposition 1

The proof of this proposition is based analysis of the repair subspaces sent by nodes in $\mathcal{A}$ to repair node x, namely $\mathcal{S}_j^x[\mathcal{A}]$ for $j \in \mathcal{A}$. In the rest of this section, since x and $\mathcal{A}$ are fixed, we may drop superscript x and parameter $\mathcal{A}$ and write $\mathcal{S}_j \triangleq \mathcal{S}_j^x[\mathcal{A}]$ for the sake of simplicity.

Before we present the proof of Proposition 1, we state and prove a few more properties of the subspaces sent by nodes in the repair process of another node.

Proposition 4. *Let $\mathcal{C}$ be a linear exact-repair regenerating code operating at an optimum point (α, β, F). The repair subspaces sent by nodes in $\mathcal{A}$ in order to repair node x are mutually linearly independent. That is, if $\sum_{j\in\mathcal{A}} \boldsymbol{v}_j = \mathbf{0}$, holds for some $\boldsymbol{v}_j \in \mathcal{S}_j$, then $\boldsymbol{v}_j = \mathbf{0}$ for every $j \in \mathcal{A}$.*

Proof of Proposition 4: We prove this by contradiction. Assume there exist vectors $\boldsymbol{v}_j \in \mathcal{S}_j$ with at least one non-zero vector (say $\boldsymbol{v}_\ell \neq \mathbf{0}$) that sum up to zero. We have

$$\boldsymbol{v}_\ell = \sum_{j\in\mathcal{A}\setminus\{\ell\}} (-\boldsymbol{v}_j) \in \sum_{j\in\mathcal{A}\setminus\{\ell\}} \mathcal{S}_j. \tag{9}$$

Since $\mathbf{0} \neq \boldsymbol{v}_\ell \in \mathcal{S}_\ell^x[\mathcal{A}]$ and $\dim(\mathcal{S}_\ell) = \beta$, there exist a subspace $\hat{S}_\ell \subset \mathcal{S}_\ell$ such that $\mathcal{S}_\ell = \hat{S}_\ell \oplus \langle \boldsymbol{v}_\ell \rangle$ and $\dim\left(\hat{\mathcal{S}}_\ell\right) \le$

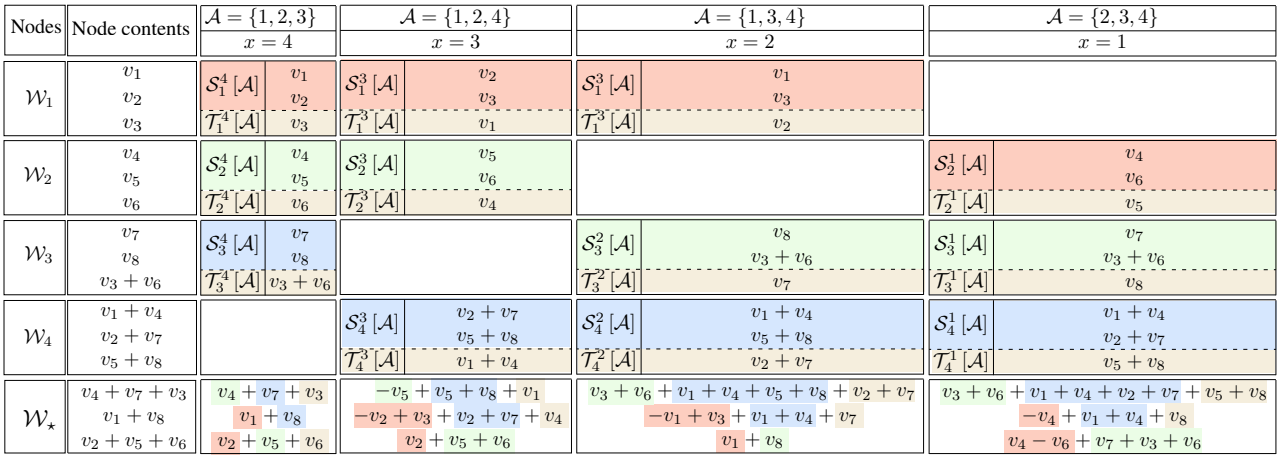

Fig. 5. Construction of the 5-th node in a $(5,3,3)$-DSS based on a known $(4,3,3)$. The first demonstrates node contents, and the next columns show $\mathcal{W}_\star$ is simultaneously well-aligned w.r.t. every $\mathcal{A}$ with $|\mathcal{A}|=3$.

$\beta-1=k-2$. Now, from (9) we have

$$\mathcal{S}_\ell=\hat{S}_\ell\oplus\langle\boldsymbol{v}_\ell\rangle\subset\hat{S}_\ell+\sum_{j\in\mathcal{A}\setminus\{\ell\}}\mathcal{S}_j. \tag{10}$$

Next, note that $|(\mathcal{A}\setminus\{\ell\})\cup\{x\}|=k$. Therefore, we have

$$\begin{aligned}\mathcal{F}&=\sum_{j\in(\mathcal{A}\setminus\{\ell\})\cup\{x\}}\mathcal{W}_j=\mathcal{W}_x+\sum_{\mathcal{A}\setminus\{\ell\}}\mathcal{W}_j\\&\subseteq\sum_{j\in\mathcal{A}}\mathcal{S}_j+\sum_{j\in\mathcal{A}\setminus\{\ell\}}\mathcal{W}_j\\&=\mathcal{S}_\ell+\sum_{j\in\mathcal{A}\setminus\{\ell\}}\mathcal{S}_j+\sum_{j\in\mathcal{A}\setminus\{\ell\}}\mathcal{W}_j\\&\overset{(a)}{\subseteq}\hat{\mathcal{S}}_\ell+\sum_{j\in\mathcal{A}\setminus\{\ell\}}\mathcal{S}_j+\sum_{\mathcal{A}\setminus\{\ell\}}\mathcal{W}_j\\&\overset{(b)}{=}\hat{\mathcal{S}}_\ell+\sum_{j\in\mathcal{A}\setminus\{\ell\}}\mathcal{W}_j\end{aligned} \tag{11}$$

where (a) is implied by (10), and (b) holds since $\mathcal{S}_j\subseteq\mathcal{W}_j$. From (11) we have

$$\begin{aligned}k^2-1=\dim(\mathcal{F})&\le\dim\left(\hat{\mathcal{S}}_\ell\right)+\sum\nolimits_{\mathcal{A}\setminus\{\ell\}}\dim(\mathcal{W}_j)\\&=(k-2)+(k-1)k=k^2-2,\end{aligned}$$

which is in infeasible. Thus the initial assumption is wrong, and repair subspaces are mutually linearly independent. ■

Proposition 5. *For a fixed pair $(\mathcal{A},x)$, define $\mathcal{S}=\sum_{i\in\mathcal{A}}\mathcal{S}_i^x[\mathcal{A}]=\sum_{i\in\mathcal{A}}\mathcal{S}_i$. Then we have $\mathcal{W}_i\setminus\mathcal{S}\neq\varnothing$ for every $i\in\mathcal{A}$.*

Proof of Prposition 5: We again prove this claim by contradiction. Suppose there exists a node $\ell\in\mathcal{A}$ such that $\mathcal{W}_\ell\subseteq\mathcal{S}$. Now, fix a node $j\in\mathcal{A}$ with $j\neq\ell$. We have

$$\begin{aligned}\mathcal{W}_\ell\subseteq\mathcal{S}=\sum_{i\in\mathcal{A}}\mathcal{S}_i&=\mathcal{S}_\ell+\mathcal{S}_j+\sum_{i\in\mathcal{A}\setminus\{j,\ell\}}\mathcal{S}_i\\&\subseteq\mathcal{S}_\ell+\mathcal{S}_j+\sum_{i\in\mathcal{A}\setminus\{j,\ell\}}\mathcal{W}_i.\end{aligned} \tag{12}$$

Next, since $|\{x\}\cup\mathcal{A}\setminus\{j\}|=k$, we have

$$\begin{aligned}\mathcal{F}&=\sum_{i\in\{x\}\cup\mathcal{A}\setminus\{j\}}\mathcal{W}_i=\mathcal{W}_x+\mathcal{W}_\ell+\sum_{i\in\mathcal{A}\setminus\{j,\ell\}}\mathcal{W}_i\\&\overset{(c)}{\subseteq}\left(\sum_{i\in\mathcal{A}}\mathcal{S}_i\right)+\left(\mathcal{S}_\ell+\mathcal{S}_j+\sum_{i\in\mathcal{A}\setminus\{j,\ell\}}\mathcal{W}_i\right)+\sum_{i\in\mathcal{A}\setminus\{j,\ell\}}\mathcal{W}_i\\&=\mathcal{S}_\ell+\mathcal{S}_j+\sum_{i\in\mathcal{A}\setminus\{j,\ell\}}(\mathcal{W}_i+\mathcal{S}_i)\overset{(d)}{=}\mathcal{S}_\ell+\mathcal{S}_j+\sum_{i\in\mathcal{A}\setminus\{j,\ell\}}\mathcal{W}_i\end{aligned}$$

where we used (12) in (c), and equality in (d) follows the fact that $\mathcal{S}_i\subseteq\mathcal{W}_i$. Therefore,

$$\begin{aligned}k^2-1=\dim(\mathcal{F})&\le\dim(\mathcal{S}_\ell)+\dim\left(\mathcal{S}_j\right)+\sum_{i\in\mathcal{A}\setminus\{j,\ell\}}\dim(\mathcal{W}_i)\\&=2(k-1)+(k-2)k=k^2-2,\end{aligned}$$

which is infeasible. This implies our initial assumption is not true, and therefore the claim of the proposition holds. ■

Now, we are ready to prove Proposition 1.

Proof of Proposition 1: Consider the repair of x by the help of $\mathcal{A}$, where each $j\in\mathcal{A}$ sends $\mathcal{S}_j=\mathcal{S}_j^x[\mathcal{A}]=\subseteq\mathcal{W}_j$. Since the code operates at $(\alpha,\beta,F)=(k,k-1,k^2-1)$, we have $\dim\left(\mathcal{S}_j^x[\mathcal{A}]\right)=k-1$. This proves part of the proposition.

Recall that $\mathcal{S}_j=\mathcal{S}_j^x[\mathcal{A}]\subseteq\mathcal{W}_j$. We have $\dim(\mathcal{W}_j)=k$, and $\dim\left(\mathcal{S}_j\right)=k-1$. Hence, $\mathcal{S}_j$ can be extended by a one-dimensional subspace $\mathcal{U}_j=\mathcal{U}_j^x[\mathcal{A}]\triangleq\left\langle u_j^x[\mathcal{A}]\right\rangle=\langle\boldsymbol{u}_j\rangle$ to span the entire space $\mathcal{W}_j$ for every $j\in\mathcal{A}$. In other words, there exist $u_j^x[\mathcal{A}]\in\mathcal{W}_j$ such that $\mathcal{W}_j=\mathcal{S}_j^x[\mathcal{A}]\oplus\left\langle u_j^x[\mathcal{A}]\right\rangle$. Note that in general there are many choices for $\left\langle u_j^x[\mathcal{A}]\right\rangle$, and we may continue the proof with *any* of them.

Next, note that there are a total of k^2 vectors stored in $\{\mathcal{W}_j:j\in\mathcal{A}\}$, and they all lie in $\mathcal{F}$ which is an (k^2-1) space. Hence, they cannot be all linearly independent. More precisely, there must be a set of vectors one from each $\mathcal{W}_j$, that sum up to zero. Since each vector in $\mathcal{W}_i$ can be written as $\xi_j\boldsymbol{u}_j+\boldsymbol{s_j}$ for some $\xi_j\in\mathbb{F}_q$, and some $\boldsymbol{s}_j\in\mathcal{S}_j$, we have

$$\exists\ \boldsymbol{s_j}\in\mathcal{S}_j,\xi_j\in\mathbb{F}_q,j\in\mathcal{A}:\ \sum_{j\in\mathcal{A}}(\xi_j\boldsymbol{u}_j+\boldsymbol{s_j})=\boldsymbol{0}. \tag{13}$$

We define $t_j^x[\mathcal{A}] \triangleq \xi_j \boldsymbol{u}_j + \boldsymbol{s_j} \in \mathcal{W}_j$ for every $j \in \mathcal{A}$. Hence, part (iv) of the proposition is immediately implied by (13). In order to prove part (ii) of the proposition, it suffices to show that $\xi_j \neq 0$ for $j \in \mathcal{A}$. First note that there is at least one non-zero ξ_j, because otherwise (13) implies existence of non-zero $\boldsymbol{s}_j \in \mathcal{S}_j$ for $j \in \mathcal{A}$ with $\sum_{j\in\mathcal{A}} \boldsymbol{s}_j = \mathbf{0}$, which is in contradiction with Proposition 4. So, let $\xi_i \neq 0$. Now if $\xi_\ell = 0$ for some $\ell \in \mathcal{A}$, we have

$$u_i = \sum_{j\in\mathcal{A}\setminus\{\ell,i\}} -\xi_i^{-1}(\xi_j u_j + \boldsymbol{s}_j) - \xi_i^{-1}\boldsymbol{s}_i - \xi_i^{-1}\boldsymbol{s}_\ell$$

which implies

$$\mathcal{U}_i \subseteq \sum_{j\in\mathcal{A}\setminus\{\ell,j\}} \mathcal{W}_j + \mathcal{S}_i + \mathcal{S}_\ell \tag{14}$$

Since $|\{x\} \cup (\mathcal{A} \setminus \{\ell\})| = k$, we have

$$\begin{aligned}
\mathcal{F} &= \mathcal{W}_x + \sum_{j\in\mathcal{A}\setminus\{\ell\}} \mathcal{W}_j = \mathcal{W}_x + \mathcal{W}_i + \sum_{i\in\mathcal{A}\setminus\{\ell,i\}} \mathcal{W}_j \\
&= \mathcal{W}_x + (\mathcal{U}_i + \mathcal{S}_i) + \sum_{i\in\mathcal{A}\setminus\{\ell,i\}} \mathcal{W}_j \\
&\overset{(e)}{\subseteq} \left(\sum_{j\in\mathcal{A}} \mathcal{S}_j\right) + \left(\sum_{j\in\mathcal{A}\setminus\{\ell,j\}} \mathcal{W}_j + \mathcal{S}_i + \mathcal{S}_\ell\right) + \sum_{j\in\mathcal{A}\setminus\{\ell,i\}} \mathcal{W}_j \\
&= \mathcal{S}_i + \mathcal{S}_\ell + \sum_{j\in\mathcal{A}\setminus\{\ell,i\}} \mathcal{W}_j, \qquad (15)
\end{aligned}$$

where (e) is implied by (14). Therefore, from (15) we have

$$k^2 - 1 = \dim(\mathcal{F}) \leq \dim\left(\mathcal{S}_i + \mathcal{S}_\ell + \sum_{j\in\mathcal{A}\setminus\{\ell,i\}} \mathcal{W}_j\right) \leq k^2 - 2.$$

This implies

$$\xi_j \neq 0, \qquad \forall j \in \mathcal{A}, \tag{16}$$

which completes the proof of part (ii) of the proposition.

Moreover, from (16), it is clear that $\mathbf{0} \neq t_j^x[\mathcal{A}] \notin \mathcal{S}_j$. Then, since $\dim(\mathcal{W}_j) = \dim\left(\mathcal{S}_j^x[\mathcal{A}]\right) + 1$, one can conclude that $\mathcal{W}_j = \mathcal{S}_j^x[\mathcal{A}] + \left\langle t_j^x[\mathcal{A}]\right\rangle$, which proves part (iii) of the proposition.

In order to show part (v), we note that $\mathcal{T}^x[\mathcal{A}] = \sum_{j\in\mathcal{A}} \mathcal{T}_j^x[\mathcal{A}] = \left\langle t_j^x[\mathcal{A}]\,; j \in A\right\rangle$ is spanned by k vectors. Moreover, part (iv) of the proposition implies that these k vectors are not linearly independent, and hence $\dim(t^x[\mathcal{A}]) \leq k-1$. Additionally, from part (iii) we have

$$\begin{aligned}
\mathcal{F} &= \sum_{j\in\mathcal{A}} \mathcal{W}_j = \sum_{j\in\mathcal{A}} \left(\mathcal{S}_j^x[\mathcal{A}] + \mathcal{T}_j^x[\mathcal{A}]\right) \\
&= \sum_{j\in\mathcal{A}} \mathcal{S}_j^x[\mathcal{A}] + \sum_{j\in\mathcal{A}} \mathcal{T}_j^x[\mathcal{A}] = \sum_{j\in\mathcal{A}} \mathcal{S}_j^x[\mathcal{A}] + \mathcal{T}^x[\mathcal{A}]
\end{aligned}$$

Hence, $\dim(\mathcal{F}) \leq \sum_{j\in\mathcal{A}} \dim\left(\mathcal{S}_j^x[\mathcal{A}]\right) + \dim(\mathcal{T}^x[\mathcal{A}])$, and so $\dim(\mathcal{T}^x[\mathcal{A}]) \geq (k^2-1) - k(k-1) = k-1$. This together with $\dim(\mathcal{T}^x[\mathcal{A}]) \leq k-1$ implies part (v) of the proposition.

Consider $\mathcal{B} = \mathcal{A} \setminus \{\ell\}$ for some $\ell \in \mathcal{A}$. Note from part (iv) of the proposition that $t_\ell^x[\mathcal{A}] = -\sum_{j\in\mathcal{B}} t_j^x[\mathcal{A}]$, which implies $\mathcal{T}_\ell^x[\mathcal{A}] \subseteq \sum_{j\in\mathcal{B}} \mathcal{T}_j^x[\mathcal{A}]$. Therefore,

$$\mathcal{T}^x[\mathcal{A}] = \sum_{j\in\mathcal{A}} \mathcal{T}_j^x[\mathcal{A}] = \mathcal{T}_\ell^x[\mathcal{A}] + \sum_{j\in\mathcal{B}} \mathcal{T}_j^x[\mathcal{A}] = \sum_{j\in\mathcal{B}} \mathcal{T}_j^x[\mathcal{A}],$$

which means $\mathcal{T}^x[\mathcal{A}]$ is spanned by any $(k-1)$ subspaces of form $\mathcal{T}_j^x[\mathcal{A}]$. This completes the proof of (vi). ■

REFERENCES

[1] A. G. Dimakis, P. B. Godfrey, Y. Wu, M. J. Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," *Information Theory, IEEE Transactions on*, vol. 56, no. 9, pp. 4539–4551, 2010.

[2] C. Tian, "Rate region of the (4, 3, 3) exact-repair regenerating codes," in *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*. IEEE, 2013, pp. 1426–1430.

[3] S. Mohajer and R. Tandon, "Exact repair for distributed storage systems: Partial characterization via new bounds," in *Information Theory and Applications Workshop (ITA), San Diego*, 2015.

[4] C. Tian, "A note on the rate region of exact-repair regenerating codes," *arXiv preprint arXiv:1503.00011*, 2015.

[5] B. Sasidharan, K. Senthoor, and P. V. Kumar, "An improved outer bound on the storage-repair-bandwidth tradeoff of exact-repair regenerating codes," in *Information Theory (ISIT), 2014 IEEE International Symposium on*. IEEE, 2014, pp. 2430–2434.

[6] S. Mohajer and R. Tandon, "New bounds on the (n, k, d) storage systems with exact repair," in *submitted to) IEEE International Symposium on Information Theory, Hong Kong*, 2015.

[7] M. Elyasi, S. Mohajer, and R. Tandon, "Linear exact repair rate region of $(k+1, k, k)$ distributed storage systems: A new approach," in *Information Theory Proceedings (ISIT), 2015 IEEE International Symposium on*. IEEE, 2015.

[8] N. Prakash and M. N. Krishnan, "The storage-repair-bandwidth trade-off of exact repair linear regenerating codes for the case $d = k = n-1$," *arXiv preprint arXiv:1501.03983*, 2015.

[9] I. M. Duursma, "Shortened regenerating codes," *arXiv preprint arXiv:1505.00178*, 2015.

[10] K. V. Rashmi, N. B. Shah, and P. V. Kumar, "Optimal exact-regenerating codes for distributed storage at the msr and mbr points via a product-matrix construction," *Information Theory, IEEE Transactions on*, vol. 57, no. 8, pp. 5227–5239, 2011.

[11] C. Suh and K. Ramchandran, "On the existence of optimal exact-repair mds codes for distributed storage," *arXiv preprint arXiv:1004.4663*, 2010.

[12] V. R. Cadambe, S. A. Jafar, and H. Maleki, "Distributed data storage with minimum storage regenerating codes-exact and functional repair are asymptotically equally efficient," *arXiv preprint arXiv:1004.4299*, 2010.

[13] C. Tian, B. Sasidharan, V. Aggarwal, V. A. Vaishampayan, and P. V. Kumar, "Layered, exact-repair regenerating codes via embedded error correction and block designs," *arXiv preprint arXiv:1408.0377*, 2014.

[14] S. Goparaju, S. E. Rouayheb, and R. Calderbank, "New codes and inner bounds for exact repair in distributed storage systems," *arXiv preprint arXiv:1402.2343*, 2014.

[15] I. Tamo, Z. Wang, and J. Bruck, "Zigzag codes: Mds array codes with optimal rebuilding," *Information Theory, IEEE Transactions on*, vol. 59, no. 3, pp. 1597–1616, 2013.

[16] D. E. Knuth, "Subspaces, subsets, and partitions," *Journal of Combinatorial Theory, Series A*, vol. 10, no. 2, pp. 178–180, 1971.